

STATEMENT OF DR. DAVID P. REED

Adj. Professor, The Media Laboratory

The Communications Futures Program

Massachusetts Institute of Technology

Weisner Building E15-492

20 Ames Street

Cambridge, Massachusetts 02139

to

Subcommittee on Telecommunications and the Internet

Committee on Energy and Commerce

U.S. House of Representatives

Washington, DC 20515-6115

17 July 2008

Mr. Chairman and Members of the Subcommittee, I thank you for the opportunity to address you on the topic of “What Your Broadband Provider Knows About Your Web Use: Deep Packet Inspection and Communications Laws and Policies.” The subject of this hearing is an important one to the country and to society as the use of the Internet becomes, more and more, a central part of every citizen's everyday life, in commerce, political expression, and culture. For the last 35 years, I have been personally involved in developing many of the key technologies of the Internet, distributed personal computing, and information sharing that we now all take for granted.

A brief summary of my main points is in order here.

First, participating in the Internet as a transport or access provider implies adherence to a set of standard technical protocols and technical practices that have been and remain essential for the proper functioning of the world-wide Internet for all its users.

Second, there is a strong distinction made in the Internet's design between information needed for transporting Internet Datagrams, and the information the Internet carries between end-point systems attached to the Internet. This distinction has a major impact on the scalability, innovation rate, and economic impact of the Internet, as well as playing an important role in ensuring privacy and safety of the users of the Internet, and limiting liability for companies that invest in providing the Internet infrastructure.

Third, technical innovations now available at very low cost in the marketplace have started to make it possible on a large scale to dig into the content of end-point to end-point messages at almost any point in the Internet transport, do selective recording and analysis of such messages, and to modify or to inject messages into the Internet that appear to be messages from a particular source, but in fact are from a third party, without the ability of the end-point systems to detect the modifications or insertions.

These technical innovations, which might be called Realtime Packet Inspection and Realtime Packet Updating, are being packaged into applications and sold as “solutions” to Internet Access Providers and Internet Transport Providers by several vendors, notably Phorm, NebuAd, Sandvine, and Ellacoya Networks, but hardly limited to those

vendors. A subset of these technologies, called Deep Packet Inspection technologies, are particularly worrisome, because they involve inspection of end-user to end-user information content, decoding, and the making of inferences about users' personal interests, private activities, etc.

In this statement, based on my expertise and direct experience as a developer and researcher for the last 35 years, a technical and marketplace-oriented discussion of these systems, their capabilities, and the uses advocated for them by their developers and by Internet transport operators, including companies such as cable, telephone and wireless carriers, that sell high-speed Internet access as part of a “Broadband” offering.

I focus my attention on the uses proposed for Deep Packet Inspection and systems supporting those uses that are being marketed to Broadband Internet Access Providers, since such providers enjoy a strong monopoly or oligopoly position in the Internet's actual deployment.

Following the discussion, I draw several conclusions that Congress may want to consider as it explores the use of these technologies.

First, that such technologies are *not at all necessary to operating the Internet or to profitable operation of an Internet operator*, and in fact that they actually violate long-agreed standards and principles that have been part of the Internet's design from the beginning, and which have led to its enormous impact and continued success.

Second, that deployment of such technologies pose *major risks to the economic success of the Internet* as a whole. They do so by normalizing non-standard and risky technical activity on the part of telecom operators who may choose to exploit captive customers, rather than transparently deliver the communications services for which their customers have paid.

Third, that protecting themselves from the negative impacts of these technologies on their private business *imposes significant additional costs on the knowledgeable customers* of the Internet transport operators and on the developers of new Internet applications, while at the same time *exploiting the unwitting and captive customers* of service providers who choose to deploy them.

My background

From the title and overview of this hearing, I understand you are interested in technical issues (such as the deployment of these technologies and their potential impact on privacy), in legal issues, and in policy-related issues. Perhaps you are also interested in their impact on innovation of Internet services, and in possible technical and legislative steps that might be taken to mitigate negative impacts on society.

Other witnesses you will hear are far more qualified than I to discuss the applicable laws and the various policy implications of these technologies. My experience and knowledge is largely in the spheres of technology, architecture and applications, based on more than 35 years of activity in computer systems, Internet communications, computer security,

and computer applications design, development, and technology strategy, both in research and industry.

Separation of concerns in the Internet Architecture

Survival of the Internet requires that Internet Access Providers and Transport Providers continue to take a proper, transparent role as participants in the Internet.

Internet Access Providers (and in particular Broadband providers offering so-called high-speed Internet *access* service) do *not* create the Internet for their customers, instead they provide *access* to the larger collective system called the Internet, of which they are a small part.

The Internet itself is the “network of networks” that results from voluntary interoperability among a wide variety of Autonomous Systems – networks that are not owned by each other, and which do not even have contractual obligations to each other in most cases. All it takes to be part of the Internet as an Autonomous System is to agree to participate according to the very simple ground rules of the Internet.¹ These ground rules are directly responsible for the remarkable growth, scalability, and resilient evolution of the Internet itself, and more importantly the growth of the Internet's utility as a backbone of commerce, information exchange, and cultural growth.

The fundamental agreement among Autonomous Systems is that they collectively

¹ The core ground rules of the Internet were laid out in the original design begun in 1975 by Vint Cerf and Bob Kahn of ARPA. I participated in that original development of, and have since written extensively about, these Internet ground rules.

provide each *host*, that is each computer that is connected to any of the many Autonomous Systems, the ability to send and receive small messages called Internet Datagrams, to any of the other hosts on any Autonomous System in the Internet. I avoid defining a whole collection of technical terms by suggesting that you view these Internet Datagrams as *envelopes* containing messages from one host to another on the Internet. The envelope is stamped on the outside with *only* four things:

- an address,
- a return address,
- a protocol identifier, and
- some marks that indicate how the message is handled as it is carried through the network.

The content of each message is held “inside the envelope.” This content is meant to be meaningful only to the sending and receiving hosts, while the envelope exterior is meant to contain all the information needed for that content to be carried from the source to the intended destination.

As a condition of participation in the Internet, each Autonomous System must agree to provide “best efforts” delivery of these Internet Datagrams (envelopes) without reading or changing their contents – that is, a sender posts an envelope with its return address and a specified destination address, and it expects that the envelope will be routed

through the network and delivered eventually to the specified address.

The concept of “outside the envelope” and “inside the envelope” is a reflection of much effort on the part of the Internet designers when the Internet was first created, and is acknowledged by many as one of the two or three reasons why

1. the Internet has scaled by many orders of magnitude over the past 30 years without a fundamental architectural change,
2. the Internet easily evolved to incorporate technological innovations in digital transport such as optical fiber switching, WiFi, 3G cellular, etc., and
3. the Internet has catalyzed the invention of a wide variety of consumer and business content distribution, communications applications and resource sharing services that range from the World Wide Web to Instant Messaging, Social Networking, business process outsourcing, etc.

It is worth thinking about these points carefully. The core idea of the Internet Datagram is a form of radical simplicity. All the Internet does is carry envelopes of a standard form – in a sense just like the post office.

Scaling: To make a faster Internet, all one need do is process the envelopes faster. To make a larger Internet, all one need to do is improve the processing units that use the address on each envelope to sort the envelope into one of the many outgoing paths from each routing point.

Technology evolution: To incorporate a new technology like optical fiber into the network, all one need do is find a way to put the bits of an Internet Datagram into a sequence of light pulses that travel down the fiber. There are numerous technical details involved in doing so, which are the province of companies like Cisco and other Internet technology providers.

Application/Service innovation: To build a new application or service, all one need do is write a program to run on standard computer servers and standard personal computer clients that communicates using a *protocol* based on Internet Datagrams. A protocol is a set of conventions or rules that specify messages that are sent inside the envelopes, in particular saying what the messages mean to the recipient, and what actions the recipient of a message should take upon the receipt of a message. Each new kind of application or service on the Internet is created by inventing a new protocol.

The Internet transport infrastructure does its job without needing to understand or to generate protocol-required messages in Application or Service protocols. Therefore, applications and services can be invented and deployed without having to negotiate consent or ask for favors from the infrastructure. The infrastructure – all of the AS's – does the same thing for *every* application: transport the envelopes.

It is this separation of concerns that is the essence of the success of the Internet.

Real-Time Packet Inspection and Real-Time Packet Updating

Because silicon computing technology has followed Moore's Law, with the size and

performance of computers improving by a factor of four every three years, the ability to process information carried in messages has improved drastically in the last 35 years. That means that today's silicon chips can, in principle, examine and process a message of a particular size about 8 *million* times more efficiently than the silicon chips at the time I began working on research computer networks in 1973.

The result of this technology evolution is that it is now quite reasonable to construct specialized computing devices that can scan tens of millions or even billions of bits of data per second passing through a network switch, running complex pattern matching and decision algorithms on each Internet Datagram during the time the Internet Datagram is received into a network switch and transmitted out over a fiber or cable to the next switch on the path between the source and destination. Since each Internet Datagram is stored in specialized buffer memory before being retransmitted, specialized devices can put put selected Datagrams aside for complex processing and modification before forwarding them on to the destination as desired.

When such devices are produced in volume, the cost per device can be made quite small. Such devices capable of monitoring, decoding, and matching Internet Datagrams are a natural result of the evolution of high performance Internet switches, but are capable of far more general operations than delivering datagrams to their intended destination.

These devices are what I call Real-Time Packet Inspection and Real-Time Packet Updating systems.

Deep Packet Inspection: Inspecting “inside the envelope”

As mentioned earlier, a core element of the Internet architecture is that the content “inside the envelope” of an Internet Datagram is not to be read or modified by any of the AS's that carry the Internet Datagram from source end-point to destination end-point.

The term *Deep Packet Inspection* was invented to describe real-time packet inspection systems that inspect and use content from “inside the envelope.” Since that content is intended only for the destination end-point, it is never necessary for AS's to inspect or analyze such content to perform their function of delivery, just as it is never necessary for the Post Office to inspect the contents of a First Class Letter in order to properly deliver the letter to its destination.

Nevertheless, a variety of companies have developed systems based on Deep Packet Inspection and have begun to market them to network transport operators and others. These systems are marketed for a variety of applications, which I will discuss below.

Often lumped into the category of Deep Packet Inspection are other techniques that involve real-time modification of the content “inside the envelope” of Internet Datagrams, and even creation of Internet Datagrams with content not requested by the source whose address appears on the “outside of the envelope”.

Though some call this modification or synthesis of Internet Datagrams “forgery” of Internet Datagrams, the legality of performing such operations involves non-technical, legal issues where I am not an expert.

Instead I will point out that the Internet Architecture, as defined by the IETF and other bodies who oversee the Internet's evolution, *neither requires nor allows* Internet Datagrams to be modified or created by AS's in this manner. I will discuss the implications and risks to end-users and the Internet as a whole of such action further below.

Thus Deep Packet Inspection goes against the separation of concerns that has been a hallmark and generator of the Internet's success.

Proposed uses of Deep Packet Inspection

When there is a technical capability of the sort we are discussing that is relatively inexpensive and quite powerful, there are many potential applications. Let me briefly list some of the potential applications where this technology appear to generate interest.

Surveillance by law enforcement and intelligence collection agencies is an application area where there is strong technical value of Deep Packet Inspection, though there is no need for updates or insertion. This application includes highly selective capture and recording of selected packets – often called “lawful intercept” by telecommunications carriers – and broad data recording and capture – often called “data mining.” Since this is typically a government function, I believe this application is out of scope for this hearing, but CALEA apparently does require that Internet operators enable the application of some forms of Deep Packet Inspection for this purpose.

Another category of application that has been marketed by vendors such as Sandvine is

“traffic management” by Internet Access Providers, which was the subject of recent hearings by the FCC, where I have testified in regard to the use of such technology by Comcast to disrupt certain categories of traffic on its Internet Access Service. That particular use involves inspection “inside the envelope” of Internet Datagrams and the technique of injecting packets that appear to be generated by the TCP protocol software as if they were sent by end-points intending to cancel the connections, with the result being that certain traffic, including BitTorrent traffic carrying large files, is disrupted severely.

Given that advertising and marketing play a large and extremely valuable part in electronic commerce using the World-Wide Web and electronic messaging, the application of such technology to analyze user behavior in order to target marketing more precisely is a hot application area.

At least two separate companies, NebuAd and Phorm, have developed systems that use Deep Packet Inspection that can be deployed within any AS to scan and analyze all Internet Datagrams, both inside the envelope and outside the envelope. The results are stored in a local database, or sent to a centralized database, recording patterns of access that are analyzed to determine each user's interests, then saved. Each of these systems also provides the ability to modify or synthesize the content of Internet Datagrams containing user-requested information from vendors and information services such as Amazon, Google, and even small business websites in order to insert advertising on the

behalf of the access provider.

It is important to note that the interception of content and modification of returned content in these systems is done under the control of the Internet Access network that installs NebuAd or Phorm. The interception and modification is beyond the control of either the consumer/user or the vendor/service who are the two primary parties. While these systems may incorporate “opt-out” provisions, privacy safeguards of the databases they construct, etc., it is important to know that their service is not a normal or accepted part of the Internet Architecture.

Another proposed use of Deep Packet Inspection technology is the scanning of traffic for undesirable, unwanted, or unlawful content. It has been proposed that Deep Packet Inspection can be used to detect unlicensed distribution of copyrighted content such as digitized movies, unwanted bulk email (spam), computer viruses, pornography, child pornography, etc. between witting and unwitting endpoints. There have been proposals that colleges, universities, and businesses, as well as Internet AS's be mandated to install such systems either by law or by legal precedents making them liable for carrying such traffic between end-points.

Finally, Deep Packet Inspection technologies are used for monitoring the performance and health of Internet operations. With such diagnostic tools, engineers can measure activity on the network, plan for facilities investments, etc. Such tools can be quite helpful in finding faults within the network and predicting areas of growth that support

AS's

This list of potential applications covers the applications of which I am current aware.

Since Deep Packet Inspection is a general-purpose capability grounded in Moore's Law and tied to the advancement of the technologies already built into the Internet switching gear being sold to customers, it will always be tempting for entrepreneurs to invent new applications for this general approach of reading, capturing, modifying, and injecting Internet Datagrams as they flow through the network.

Risks associated with Deep Packet Inspection

As noted earlier, a very useful way to think about the how the Internet is structured is to imagine Internet Datagrams as packages or envelopes carried by a sequence of third party delivery services from an end-point computer to another end-point computer. In this analogy, the Autonomous Systems are like individual package delivery services, such as UPS, FedEx, DHL, Yellow Trucking, etc. On the “outside” of packages is a set of labels that are intended for use in forwarding the package on to its destination. A carrier that picks up a package may transfer the package from one carrier to another, choosing the path best suited for timely delivery of the package to its destination.

In this analogy, Deep Packet Inspection technologies can be accurately thought of as devices that are placed in trucks, airplanes, warehouses, etc. of the various forwarding services that very quickly and efficiently examine the contents of the packages (perhaps by X-ray, by actually opening the package and taking pictures of its content, ...), record

the results of that examination in a database held by a third party, and analyze all of the information captured using statistical methods. This information is then used to select particular packages for special handling, discarding, or re-routing, to change, delete, or insert contents into the packages, and to create packages that *appear to be* from a particular source, but which are in fact generated by within the network itself.

A useful example in this analogy would be if all of the packages you ship and receive through an independent shipping agent (such as Mailboxes, Etc.) were scanned, and the contents used to understand your buying habits, and further that the shipping agent had a contract with various companies to insert or replace the contents of your packages with “improved” contents.

I suspect that there may be some users who would be delighted to receive items they did not request in their packages, and that merchants might be happy to find that the computer that they ship to a customer is magically “improved” or replaced by an upgraded model along the way.

However, the normal understanding in dealing with shipping agents is that the contents of packages are not to be examined, studied, reported to third parties, replaced or modified according to the desires of third parties.

There is another problem, however, that to me is more problematic. That is that unlike the shipping agent example above, the Internet is based on end-to-end protocols that are more complex than simply the delivery of packages. In these protocols, the contents of

packages contain requests for remote end-point service systems to perform actions on the user's behalf, which generate responses and then more requests. As an example, consider a user who coordinates his or her finances among of number of banks and brokers via protocols carried in these packages. He or she might first send a deposit to one account, then request a transfer to another bank once the deposit is confirmed, and then send instructions about investing the money to the second bank. This sequence of steps is called a protocol.

What Deep Packet Inspection technologies attempt to do is to *second-guess the intent* of the end-users of these services (the investor and the bank), to draw inferences about the intent of those protocols and to modify (hopefully safely) the packages without causing harm to the protocol transactions.

The source of the problem is that *vendors of Deep Packet Inspection systems cannot presume to understand, merely by looking at the contents of packages what they actually mean or intend to happen at the source and endpoint.*

This is the real risk: an service or technology *unnecessary to the correct functioning of the Internet* is introduced at a place where it cannot function correctly because it does not know the endpoints' intent, yet it operates invisibly and violates rules of behavior that the end-users and end-point businesses depend to work in a specific way.

As a simple example, I cannot send email from many hotels, because of a Deep Packet Inspection technology deployed in many hotels' Internet Access service. That service

(intended to block spammers who might operate from hotel rooms) intercepts my packages intended for my email server, and responds, pretending to be any and all destination email servers, offering to accept my email messages on behalf of the recipients, which it will then scan for evidence of viruses and spam. In my case, I use a special secure, encrypted email delivery service that is more secure than most, so my mail sending software recognizes the deception and refuses to deliver the mail to the deceptive provider that requires me to send my mail “in the clear” so it can be scanned.

Hotel providers claim that they are “doing a service” by blocking spam, but in doing so they reduce my own personal security, both by requiring that my mail be sent in the clear, and by introducing the risk that my mail will be scanned and modified by an interceptor I cannot easily avoid. Some hotel providers even claim that they are legally *mandated by liability law* to inspect my email that originates through the hotel system.

In addition, if my email requests happen to involve the transmission of messages that the operator *deems to be spam*, my message, which may be quite important to my business, will be blocked without my knowledge or any possibility to appeal the erroneous inference.

That example, though a simple example, captures the risks that I want to highlight:

- Systems based on Deep Packet Inspection work by drawing inferences from packet contents that are not intended to be understood by anyone other than the destination host. Deep Packet Inspection systems *cannot* reliably determine the

intent or meaning of those Internet Datagrams.

- Deep Packet Inspection systems work by deliberately interfering with end-to-end communications, but by definition attempt to deceive the endpoint systems about what the original Internet Datagrams contain. The endpoints cannot tell if such systems have either captured their content information, or modified or created information that was not sent or intended by the author of the Internet Datagram.
- Deep Packet Inspection systems cannot be made reliable, either in their inference or in their actions.

Impacts on Users and Services Built on the Internet

In order to block interception and modification of the contents of their Internet

Datagrams, end-point hosts can take steps such as encrypting contents of packets, using digital signatures, and choosing providers that vow not to scan or modify packets.

Besides raising the cost of using the Internet for existing and new applications, there are three problems with this.

First, existing applications have been designed with the expectation that Deep Packet Inspection is not a legitimate activity by a service provider.

Second, there is only one Internet, which consists of many Autonomous Systems.

Choosing a different point of connection cannot, given the nature of the Internet, ensure that all users one might want to send Internet Datagrams to have successfully chosen

providers that have not deployed Deep Packet Inspection systems that scan or modify Internet Datagrams. Thus, consumer choice is not an option. Since the risks of incorrect operation of Deep Packet Inspection can disrupt critical protocols (including protocols yet to be deployed or invented), mere consumer choice may not be enough to fix the problem.

Third, encryption from end-to-end, while a potential solution, has public policy implications. This committee and Congress have gone through those issues many times. I personally would like to see all communications activities fully protected by strong encryption, but I fear that reaching that point will encounter many obstacles. If the primary problem the encryption is to deal with is an unnecessary technology such as Deep Packet Inspection, a simpler solution would be to bar the use of Deep Packet Inspection systems.